arXiv:1604.01996v1 [stat.ME] 7 Apr 2016

CopulaDTA: An R Package for Copula Based Bivariate Beta-Binomial Models for Diagnostic Test Accuracy Studies in a Bayesian Framework

Victoria N Nyaga
Institute of Public Health,
Hasselt University

Marc Arbyn
Institute of Public Health

Marc Aerts
Hasselt University

Abstract

The current statistical procedures implemented in statistical software packages for pooling of diagnostic test accuracy data include hSROC regression (Rutter and Gatsonis 2001) and the bivariate random-effects meta-analysis model (BRMA) (Reitsma *et al.* (2005), Arends *et al.* (2008), Chu and Cole (2006), Riley *et al.* (2007b)). However, these models do not report the overall mean but rather the mean for a central study with random-effect equal to zero and have difficulties estimating the correlation between sensitivity and specificity when the number of studies in the meta-analysis is small and/or when the between-study variance is relatively large (Riley *et al.* 2007a).

This tutorial on advanced statistical methods for meta-analysis of diagnostic accuracy studies discusses and demonstrates Bayesian modeling using **CopulaDTA** (Nyaga 2015) package in R (R Core Team 2015) to fit different models to obtain the meta-analytic parameter estimates. The focus is on the joint modelling of sensitivity and specificity using copula based bivariate beta distribution. Essentially, we extend the work of Nikoloulopoulos (2015a) by: **i**) presenting the Bayesian approach which offers flexibility and ability to perform complex statistical modelling even with small data sets and **ii**) including covariate information, and **iii**) providing an easy to use code. The statistical methods are illustrated by re-analysing data of two published meta-analyses.

Modelling sensitivity and specificity using the bivariate beta distribution provides marginal as well as study-specific parameter estimates as opposed to using bivariate normal distribution (e.g., in BRMA) which only yields study-specific parameter estimates. Moreover, copula based models offer greater flexibility in modelling different correlation structures in contrast to the normal distribution which allows for only one correlation structure.

Keywords: diagnostic test accuracy, meta-analysis, Bayesian, random-effects, copula, R.

1. Introduction

In a systematic review of diagnostic test accuracy, the statistical analysis section aims at estimating the average (across studies) sensitivity and specificity of a test and the variability thereof, among other measures. There tends to be a negative correlation between sensitivity and specificity, which postulates the need for correlated data models. The analysis is statistically challenging because the user **i**) deals with two summary statistics, **ii**) has to account for correlation between sensitivity and specificity, **iii**) has to account for heterogeneity in sensitivity and specificity across the studies and **iv**) should be allowed to incorporate covariates.

Currently, the HSROC regression (Rutter and Gatsonis 2001) or the bivariate random-effects meta-analysis model (BRMA) (Reitsma *et al.* (2005), Arends *et al.* (2008), Chu and Cole (2006)) are recommended for pooling of diagnostic test accuracy data. These models fit a bivariate normal distribution which allows for only one correlation structure to the logit transformed sensitivity and specificity. The resulting distribution has no closed form and therefore the mean sensitivity and specificity is only estimated after numerical integration or other approximation methods, an extra step which is rarely taken.

Within the maximum likelihood estimation methods, the BRMA and HSROC models have difficulties in convergence and estimating the correlation parameter when the number of studies in the meta-analysis are small and/or when the between-study variances are relatively large (Takwoingi *et al.* 2015). When the correlation is close to the boundary of its parameter space, the between study variance estimates from the BRMA are upwardly biased as they are inflated to compensate for the range restriction on the correlation parameter (Riley *et al.* 2007a). According to Riley *et al.* (2007b) this occurs because the maximum likelihood estimator truncates the between-study covariance matrix on the boundary of its parameter space, and this often occurs when the within-study variation is relatively large or the number of studies is small.

The BRMA and HSROC assume that the transformed data is approximately normal with constant variance, however for sensitivity and specificity and proportions in generals, the mean and variance depend on the underlying probability. Therefore, any factor affecting the probability will change the mean and the variance. This implies that the in models where the predictors affect the mean but assume a constant variance will not be adequate.

Joint modelling of study specific sensitivity and specificity using existing or copula based bivariate beta distributions overcomes the above mentioned difficulties. Since both sensitivity and specificity take values in the interval space (0, 1), it is a more natural choice to use a beta distribution to describe their distribution across studies, without the need for any transformation. The beta distribution is conjugate to the binomial distribution and therefore it is easy to integrate out the random-effects analytically giving rise to the beta-binomial marginal distributions. Moreover no further integration is needed to obtain the meta-analytically pooled sensitivity and specificity. Previously, Cong *et al.* (2007) fitted separate beta-binomial models to the number of true positives and the number of false positives. While the model ignores correlation between sensitivity and specificity, Cong *et al.* (2007) reported that the model estimates are comparable to those from the SROC model (Moses *et al.* 1993), the predecessor of the HSROC model.

According to Riley (2009), ignoring the correlation would have negligible influence on the meta-analysis results when the within-study variability is large relative to the between-study variability. By utilising the correlation, we allow borrowing strength across sensitivities and specificities resulting in smaller standard errors. The use of copula based mixed models

within the frequentist framework for meta-analysis of diagnostic test accuracy was recently introduced by Nikoloulopoulos (2015a) who evaluated the joint density numerically.

This tutorial, presents and demonstrates hierarchical mixed models for meta-analysis of diagnostic accuracy studies. In the first level of the hierarchy, given sensitivity and specificity for each study, two binomial distributions are used to describe the variation in the number of true positives and true negatives among the diseased and healthy individuals, respectively. In the second level, we model the unobserved sensitivities and specificities using a bivariate distribution. While hierarchical models are used, the focus of meta-analysis is on the pooled average across studies and rarely on a given study estimate.

The methods are demonstrated using datasets from two previously published meta-analyses: **a**) on diagnostic accuracy of telomerase in urine as a tumour marker for the diagnosis of primary bladder cancer from Glas *et al.* (2003) previously used by Riley *et al.* (2007b) and Nikoloulopoulos (2015a) since it is a problematic dataset that has convergence issues caused by the correlation parameter being estimated to be -1 and has no covariate and **b**) on the comparison of the sensitivity and specificity of human papillomavirus testing (using the HC2 assay) versus repeat cytology to triage women with minor cytological cervical lesions to detect underlying cervical precancer from Arbyn *et al.* (2013). the second dataset is used to demonstrate meta-regression with one covariate which can be naturally extended to include several covariates.

The layout of this tutorial is as follows: in Section 2 introduces the concept of copula theory and different bivariate distributions for sensitivity and specificity. The software implementation and model selection in a Bayesian framework is discussed in Section 3. The two aforementioned datasets are introduced in Section 4. Application of software, code examples and the results of the models fitted to the data are presented in Section 5 and 6. A brief discussion is found in Section 7 and a conclusion in Section 8.

2. Statistical methods for meta-analysis

2.1. Definition of copula function

A bivariate copula function describes the dependence structure between two random variables. Two random variables X_1 and X_2 are joined by a copula function C if their joint cumulative distribution function can be written as

$$F(x_1,\ x_2) = C(F_1(x_1),\ F_2(x_2)),\ \ -\infty\ \leq\ x_1,\ x_2\ \leq\ \ +\infty. \tag{1}$$

According to the theorem of Sklar (1959), there exists for every bivariate (multivariate in extension) distribution a copula representation C which is unique for continuous random variables. If the joint cumulative distribution function and the two marginals are known, then the copula function can be written as

$$C(u,\ v) = F(F_1^{-1}(u),\ F_2^{-1}(v)),\ 0 \leq\ u,\ v\ \leq\ 1. \tag{2}$$

A 2-dimensional copula is in fact simply a 2-dimensional cumulative function restricted to the unit square with standard uniform marginals. A comprehensive overview of copulas and

their mathematical properties can be found in Nelsen (2006). To obtain the joint probability density, the joint cumulative distribution in Equation 1 should be differentiated to yield

$$f(x_1,\ x_2) = f_1(x_1)\ f_2(x_2)\ c(F_1(x_1),\ F_2(x_2)), \tag{3}$$

where f_1 and f_2 denote the marginal density functions and c the copula density function corresponding to the copula cumulative distribution function C. Therefore from Equation 3, a bivariate probability density can be expressed using the marginal and the copula density, given that the copula function is absolutely continuous and twice differentiable.

When the functional form of the marginal and the joint densities are known, the copula density can be derived as follows

$$c(F_1(x_1),\ F_2(x_2)) = \frac{f(x_1,\ x_2)}{f_1(x_1)\ f_2(x_2)}. \tag{4}$$

While our interest does not lie in finding the copula function, Equation 3 and 4 serve to show how one can move from the copula function to the bivariate density or vice-versa, given that the marginal densities are known. The decompositions allow for constructions of other and possible better models for the variables than would be possible if we limited ourselves to only existing standard bivariate distributions.

We finish this section by mentioning an important implication when Sklar's theorem is extended to a meta-regression setting with covariates. According to Patton (2006), it is important that the conditioning variable remains the same for both marginal distributions and the copula, as otherwise the joint distribution might not be properly defined. This implies that covariate information should be introduced in both the marginals and the association parameters of the model.

2.2. The hierarchical model

Since there are two sources of heterogeneity in the data, the within- and between-study variability, the parameters involved in a meta-analysis of diagnostic accuracy studies vary at two levels. For each study i, $i = 1, \ldots, n$, let $Y_i = (Y_{i1},\ Y_{i2})$ denote the true positives and true negatives, $N_i = (N_{i1},\ N_{i2})$ the diseased and healthy individuals respectively, and $\pi_i = (\pi_{i1},\ \pi_{i2})$ represent the 'unobserved' sensitivity and specificity respectively.

Given study-specific sensitivity and specificity, two separate binomial distributions describe the distribution of true positives and true negatives among the diseased and the healthy individuals as follows

$$Y_{ij} \mid \pi_{ij},\ \mathbf{x}_i\ \sim\ bin(\pi_{ij},\ N_{ij}), i\ =\ 1,\ \ldots\ n,\ j\ =\ 1,\ 2, \tag{5}$$

where $\mathbf{x}_i$ generically denotes one or more covariates, possibly affecting π_{ij}. Equation 5 forms the higher level of the hierarchy and models the within-study variability. The second level of the hierarchy aims to model the between study variability of sensitivity and specificity while accounting for the inherent negative correlation thereof, with a bivariate distribution as follows

$$\begin{pmatrix} g(\pi_{i1}) \\ g(\pi_{i2}) \end{pmatrix} \sim f(g(\pi_{i1}),\ g(\pi_{i2}))\ =\ f(g(\pi_{i1}))\ f(g(\pi_{i2}))\ c(F_1(g(\pi_{i1})),\ F_2(g(\pi_{i2}))), \tag{6}$$

where $g(.)$ denotes a transformation that might be used to modify the (0, 1) range to the whole real line. While it is critical to ensure that the studies included in the meta-analysis satisfy the specified entry criterion, there are study specific characteristics like different test thresholds and other unobserved differences that give rise to the second source of variability, the between-study variability. It is indeed the difference in the test thresholds between the studies that gives rise to the correlation between sensitivity and specificity. Including study level covariates allows us to model part of the between-study variability. The covariate information can and should (Patton 2006) be used to model the mean as well as the correlation between sensitivity and specificity.

In the next section we give more details on different bivariate distributions $f(g(\pi_{i1}),\ g(\pi_{i2}))$ constructed using the logit or identity link function $g(.)$, different marginal densities and/or different copula densities c. We discuss their implications and demonstrate their application in meta-analysis of diagnostic accuracy studies. An overview of suitable parametric families of copula for mixed models for diagnostic test accuracy studies was recently given by Nikoloulopoulos (2015a). Here, we consider five copula functions which can be plugged in Equation 3 to model negative correlation.

Bivariate Gaussian copula

Given the density and the distribution function of the univariate and bivariate standard normal distribution with correlation parameter $\rho \in (-1, 1)$, the bivariate Gaussian copula function and density is expressed (Meyer 2013) as

$$\begin{aligned} C(u,\ v,\ \rho) &= \Phi_2(\Phi^{-1}(u),\ \Phi^{-1}(v),\ \rho), \\ c(u,\ v,\ \rho) &= \frac{1}{\sqrt{1\ -\ \rho^2}}\ exp\left(\frac{2\ \rho\ \Phi^{-1}(u)\ \Phi^{-1}(v) - \rho^2\ (\Phi^{-1}(u)^2 + \Phi^{-1}(v)^2)}{2\ (1-\rho^2)}\right). \end{aligned} \tag{7}$$

The logit transformation is often used in binary logistic regression to relate the probability of "success" (coded as 1, failure as 0) of the binary response variable with the linear predictor model that theoretically can take values over the whole real line. In diagnostic test accuracy studies, the 'unobserved' sensitivities and specificities can range from 0 to 1 whereas their logits $= log(\frac{\pi_{ij}}{1\ -\ \pi_{ij}})$ can take any real value allowing to use the normal distribution as follows

$$logit(\pi_{ij})\ \sim\ N(\mu_j,\ \sigma_j)\ <=>\ logit(\pi_{ij})\ =\ \mu_j\ +\ \varepsilon_{ij}, \tag{8}$$

where, μ_j is a vector of the mean sensitivity and specificity for a study with zero random effects, and ε_i is a vector of random effects associated with study i. Now u is the normal distribution function of $logit(\pi_{i1})$ with parameters μ_1 and σ_1, v is the normal distribution function of $logit(\pi_{i2})$ with parameters μ_2 and σ_2, Φ_2 is the distribution function of a bivariate standard normal distribution with correlation parameter $\rho \in (-1,\ 1)$ and Φ^{-1} is the quantile of the standard normal distribution. In terms of ρ, Kendall's tau is expressed as $(\frac{2}{\pi})\arcsin(\rho)$.

With simple algebra the copula density in Equation 7 with normal marginal distributions simplifies to

$$c(u,\ v,\ \rho) = \frac{1}{\sqrt{1-\rho^2}}\ exp\left(\frac{1}{2\ (1\ -\ \rho^2)}\left(\frac{2\ \rho\ (x-\mu_1)\ (y-\mu_2)}{\sigma_1\ \sigma_2} - \rho^2\left(\frac{x\ -\ {\mu_1}^2}{\sigma_1} + \frac{y\ -\ {\mu_2}^2}{\sigma_2}\right)\right)\right). \tag{9}$$

The product of the copula density in Equation 9, the normal marginal of $logit(\pi_{i1})$ and $logit(\pi_{i2})$ in Equation 8 form a bivariate normal distribution which characterize the model by Reitsma *et al.* (2005), Arends *et al.* (2008), Chu and Cole (2006), and Riley *et al.* (2007a), the so-called bivariate random-effects meta-analysis (BRMA) model, recommended as the appropriate method for meta-analysis of diagnostic accuracy studies. Study level covariate information explaining heterogeneity is introduced through the parameters of the marginal and the copula as follows

$$\boldsymbol{\mu}_j = \mathbf{X}_j \mathbf{B}_j^{\top}. \tag{10}$$

$\mathbf{X}_j$ is a $n \times p$ matrix containing the covariates values for the mean sensitivity($j = 1$) and specificity($j = 2$). For simplicity, assume that $\mathbf{X}_1 = \mathbf{X}_2 = \mathbf{X}$. $\mathbf{B}_j^{\top}$ is a $p \times 1$ vector of regression parameters, and p is the number of parameters. By inverting the logit functions in Equation 8, we obtain

$$\pi_{ij} = logit^{-1}(\mu_j + \varepsilon_{ij}). \tag{11}$$

Therefore, the meta-analytic sensitivity and specificity obtained by averaging over the random study effect, is given by, for $j = 1, 2$

$$\mathsf{E}(\pi_j) \;=\; \mathsf{E}(logit^{-1}(\mu_j \;+\; \varepsilon_{ij})) \;=\; \int_{-\infty}^{\infty} logit^{-1}(\mu_j \;+\; \varepsilon_{ij}) f(\varepsilon_{ij},\; \sigma_j)\; d\varepsilon_{ij}, \tag{12}$$

assuming that $\sigma_1^2 > 0$ and $\sigma_2^2 > 0$. The integration in Equation 12 has no analytical expression and therefore needs to be numerically approximated and the standard are not easily available. Using MCMC simulation in the Bayesian framework the meta-analytic estimates can be easily computed as well as a standard error estimate and a credible intervals $\mathsf{E}(\pi_j)$ with minimum effort by generating predictions of the fitted bivariate normal distribution.

In the frequentist framework, it is more convenient however to use numerical averaging by sampling a large number M of random-effects $\hat{\varepsilon}_{ij}$ from the fitted distribution and to estimate the meta-analytic sensitivity and specificity by (Molenberghs and Verbeke 2005), for $j = 1, 2$

$$\hat{\mathsf{E}}(\pi_j) = \frac{1}{M} \sum_{i \,=\, 1}^{M} logit^{-1}(\hat{\mu}_j + \hat{\varepsilon}_{ij}). \tag{13}$$

However, inference is not straightforward in the frequentist framework since the standard errors are not available. When $\varepsilon_{ij} = 0$, then

$$\mathsf{E}(\pi_j \mid \varepsilon_{ij} = 0) \;=\; logit^{-1}(\mu_j). \tag{14}$$

Inference for $\mathsf{E}(\pi_j \mid \varepsilon_{ij} \;=\; 0)$, as expressed in Equation 14, can be done in both Bayesian and frequentist framework. The equation represents the mean sensitivity and specificity for a "central" study with $\varepsilon_{ij} \;=\; 0$. Researchers often seem to confuse $\mathsf{E}(\pi_j \mid \varepsilon_{ij} \;=\; 0)$ with $\mathsf{E}(\pi_j)$ but due to the non-linear logit transformations, they are clearly not the same parameter.

With the identity link function, no transformation on study-specific sensitivity and specificity is performed. A natural choice for u and v would be beta distribution functions with parameters $(\alpha_1,\; \beta_1)$ and $(\alpha_2,\; \beta_2)$ respectively. Since $\pi_{ij} \;\sim\; beta(\alpha_j,\; \beta_j)$, the meta-analytic sensitivity and specificity are analytically solved as follows

$$\mathsf{E}(\pi_j) = \frac{\alpha_j}{\alpha_j + \beta_j}, \tag{15}$$

After reparameterising the beta distributions using the mean ($\mu_j = \frac{\alpha_j}{\alpha_j + \beta_j}$) and certainty ($\psi_j = \alpha_j + \beta_j$) or dispersion ($\varphi_j = \frac{1}{1 + \alpha_j + \beta_j}$) parameters different link functions introduce covariate information to the mean, certainty/dispersion and association (ρ) parameters. A typical model parameterisation is

$$\begin{aligned}
\boldsymbol{\mu}_j &= logit^{-1}(\mathbf{X}\mathbf{B}_j^\top), \\
\boldsymbol{\psi}_j &= g(\mathbf{W}\mathbf{C}_j^\top), \\
\boldsymbol{\alpha}_j &= \boldsymbol{\mu}_j \circ \boldsymbol{\psi}_j, \\
\boldsymbol{\beta}_j &= (\mathbf{1} - \boldsymbol{\mu}_j) \circ \boldsymbol{\psi}_j, \\
\boldsymbol{\rho} &= tanh(\mathbf{Z}\mathbf{D}_j^\top) = \frac{exp(2 \times \mathbf{Z}\mathbf{D}_j^\top) - 1}{exp(2 \times \mathbf{Z}\mathbf{D}_j^\top) + 1}.
\end{aligned} \tag{16}$$

X, W and **Z** are a $n \times p$ matrices containing the covariates values for the mean, dispersion and correlation which we will assume has similar information and denoted by **X** for simplicity purpose, p is the number of parameters, $\mathbf{B}_j^\top$, $\mathbf{V}_j^\top$ and $\mathbf{D}_j^\top$ are a $p \times 1$ vectors of regression parameters relating covariates to the mean, variance and correlation respectively. $g(.)$ is the log link to mapping $\mathbf{X}\mathbf{C}_j^\top$ to the positive real number line and $\circ$ is the Hadamard product.

Frank copula

This flexible copula in the so-called family of Archimedean copulas was introduced by Frank (1979). The functional form of the copula and the density which is plugged in Equation 3 is given by;

$$\begin{aligned}
C(F(\pi_{i1}),\ F(\pi_{i2}), \theta) &= -\frac{1}{\theta}\ log\left[1 + \frac{(e^{-\theta\ F(\pi_{i1})} - 1)(e^{-\theta\ F(\pi_{i2})} - 1)}{e^{-\theta} - 1}\right], \\
c(F(\pi_{i1}),\ F(\pi_{i2}), \theta) &= \frac{\theta\ (1 - e^{-\theta})\ e^{-\theta\ (F(\pi_{i1})\ +\ F(\pi_{i2}))}}{[1 - e^{-\theta} - (1 - e^{-\theta\ F(\pi_{i1})})\ (1 - e^{-\theta\ F(\pi_{i2})})]^2}.
\end{aligned} \tag{17}$$

Since $\theta \in \mathbb{R}$, both positive and negative correlation can be modelled, making this one of the more comprehensive copulas. When θ is 0, sensitivity and specificity are independent. For $\theta > 0$, sensitivity and specificity exhibit positive quadrant dependence and negative quadrant dependence when $\theta < 0$. The Spearman correlation ρ_s and Kendall's tau τ_k can be expressed in terms of θ as

$$\begin{aligned}
\rho_s &= 1 - 12\ \frac{D_2(-\theta)\ -\ D_1(-\theta)}{\theta}, \\
\tau_k &= 1 + 4\ \frac{D_1(\theta)\ -\ 1}{\theta},
\end{aligned} \tag{18}$$

where $D_j(\delta)$ is the Debye function defined as

$$D_j(\delta) = \frac{j}{\delta^j} \int_\theta^\delta \frac{t^j}{exp(t)\ -\ 1}\ dt,\ j\ = 1,\ 2. \tag{19}$$

Covariate information is introduced in a similar manner as Equation 16. The identity link is used for the association parameter θ.

Farlie-Gumbel-Morgenstern copula (FGM)

This popular copula studied by Farlie (1960), Gumbel (1960) and Morgenstern (1956) is defined as

$$\begin{aligned} C(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) &= F(\pi_{i1})\ F(\pi_{i2})[1\ +\ \theta\ (1-F(\pi_{i1}))\ (1\ -\ F(\pi_{i2}))], \\ c(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) &= [1\ +\ \theta\ (2\ F(\pi_{i1})\ -\ 1)\ (2\ F(\pi_{i2})\ -\ 1)]. \end{aligned} \tag{20}$$

Because $\theta \in (-1, 1)$, the Spearman correlation and Kendall's tau are expressed in terms of θ as $\theta/3$ and $2\theta/9$ respectively, making this copula only appropriate for data with weak dependence since $|\rho_s| \leq 1/3$. In a similar manner as in Equation 16 the logit link, log/identity link and Fisher's z transformation can be used to introduce covariate information in modelling the mean, dispersion and association parameter.

Clayton copula

The Clayton copula function and density by Clayton (1978) is defined as

$$\begin{aligned} &C(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) = [F(\pi_{i1})^{-\theta}\ +\ F(\pi_{i2})^{-\theta}\ -\ 1]^{\frac{-1}{\theta}}, \\ &c(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) = (1\ +\ \theta)\ F(\pi_{i1})^{-(1\ +\ \theta)}\ F(\pi_{i2})^{-(1+\theta)}\ [F(\pi_{i1})^{-\theta} + F(\pi_{i2})^{-\theta}\ -\ 1]^{\frac{-(2\ \theta\ +\ 1)}{\theta}}. \end{aligned} \tag{21}$$

Since $\theta \in (0,\ \infty)$, the Clayton copula typically models positive dependence; Kendall's tau equals $\theta/(\theta\ +\ 2)$. However, the copula function can be rotated by 90° or 270° to model negative dependence. The distribution and density functions following such rotations are given by

$$\begin{aligned} C_{90}(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) &= F(\pi_{i2})\ -\ C(1\ -\ F(\pi_{i1}),\ F(\pi_{i2}),\ \theta), \\ c_{90}(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) &= (1\ +\ \theta)(1\ -\ F(\pi_{i1}))^{-(1\ +\ \theta)}\ F(\pi_{i2})^{-(1\ +\ \theta)}\ [(1-F(\pi_{i1}))^{-\theta} \\ &\quad +\ F(\pi_{i2})^{-\theta} - 1]^{\frac{-(2\ \theta\ +\ 1)}{\theta}}, \end{aligned} \tag{22}$$

and

$$\begin{aligned} C_{270}(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) &= F(\pi_{i1}) - C(F(\pi_{i1}),\ 1\ -\ F(\pi_{i2}), \theta), \\ c_{270}(F(\pi_{i1}),\ F(\pi_{i2}),\ \theta) &= (1\ +\ \theta)\ F(\pi_{i1})^{-(1\ +\ \theta)}\ (1\ -\ F(\pi_{i2}))^{-(1\ +\ \theta)}\ [F(\pi_{i1})^{-\theta} \\ &\quad +(1\ -\ F(\pi_{i2}))^{-\theta}\ -\ 1]^{\frac{-(2\ \theta\ +\ 1)}{\theta}}. \end{aligned} \tag{23}$$

The logit, log/identity and log/identity links can be used to introduce covariate information in modelling the mean (μ_j), certainty (ψ_j)/dispersion (φ_j) and association (θ) parameters respectively in the same way as in Equation 16.

Of course other copula functions that allow for negative association can be chosen. It is also an option to use known bivariate beta distributions. However, it is not always straightforward and analytically attractive to derive the corresponding copula function for all bivariate distributions. The use of existing bivariate beta distributions in meta-analysis of diagnostic accuracy studies has been limited because these densities model positive association (e.g., Libby and Novick (1982), Olkin and Liu (2003)), or both positive and negative association but over a restricted range(e.g., Sarmanov (1966)).

3. Inference Framework and Software

Within the Bayesian framework, the analyst updates a prior opinion/information of a parameter based on the observed data whereas in the frequentist framework, the analyst investigates the behaviour of the parameter estimates in hypothetical repeated samples from a certain population. Due to its flexibility and use of MCMC simulations, complex modelling can often be implemented more easily within the Bayesian framework. By manipulating the prior distributions, Bayesian inference can circumvent identifiability problems whereas numerical approximation algorithms in frequentist inference without prior distributions can become stuck caused by identifiability problems. However, Bayesian methods typically require statistical expertise and patience because the MCMC simulations are computationally intensive. In contrast, most frequentist methods have been wrapped up in standard 'procedures' that require less statistical knowledge and programming skills. Moreover frequentist methods are optimized with maximum likelihood estimation (MLE) that have much shorter run-times as opposed to MCMC simulations. **CopulaREMADA** (Nikoloulopoulos 2015b) is such a MLE based R package.

3.1. The CopulaDTA package

The **CopulaDTA** package is an extension of **rstan** (Stan 2015), the R interface to Stan (Carpenter *et al.*) for diagnostic test accuracy data. Stan is a probabilistic programming language which has implemented Hamilton Monte Carlo(MHC) and uses No-U-Turn sampler (NUTS) (Hoffman and Gelman). The package facilitates easy application of complex models and their visualization within the Bayesian framework with much shorter run-times.

JAGS (Plummer 2003) is an alternative extensible general purpose sampling engine to Stan. Extending JAGS requires knowledge of C++ to assemble a dynamic link library(DLL) module. From experience, configuring and building the module is a daunting and tedious task especially in the Windows operation system. The above short-comings coupled with the fact that Stan tends to converge with fewer iterations even from bad initial values than JAGS made us prefer the Stan MCMC sampling engine.

The **CopulaDTA** package is available via the Comprehensive R Archive Network (CRAN) at `http://CRAN.R-project.org/package=CopulaDTA`. With a working internet connection, the **CopulaDTA** package is installed and loaded in R with the following commands

```
R> install.packages("CopulaDTA", dependencies = TRUE)

R> library(CopulaDTA)
```

The **CopulaDTA** package provide functions to fit bivariate beta-binomial distributions constructed as a product of two beta marginal distributions and copula densities discussed in Section 2. The package also provides forest plots for a model with categorical covariates or with intercept only. Given the chosen copula function, a beta-binomial distribution is assembled up by the `cdtamodel` function which returns a `cdtamodel` object. The main function `fit` takes the `cdtamodel` object and fits the model to the given dataset and returns a `cdtafit` object for which `print`, `summary` and `plot` methods are provided for.

3.2. Model diagnostics

To assess model convergence, mixing and stationarity of the chains, it is necessary to check the potential scale reduction factor $\hat{R}$, effective sample size (ESS), MCMC error and trace plots of the parameters. When all the chains reach the target posterior distribution, the estimated posterior variance is expected to be close to the within chain variance such that the ratio of the two, $\hat{R}$ is close to 1 indicating that the chains are stable, properly mixed and likely to have reached the target distribution. A large $\hat{R}$ indicates poor mixing and that more iterations are needed. Effective sample size indicates how much information one actually has about a certain parameter. When the samples are auto correlated, less information from the posterior distribution of our parameters is expected than would be if the samples were independent. ESS close to the total post-warm-up iterations is an indication of less autocorrelation and good mixing of the chains. Simulations with higher ESS have lower standard errors and more stable estimates. Since the posterior distribution is simulated there is a chance that the approximation is off by some amount; the Monte Carlo (MCMC) error. MCMC error close to 0 indicates that one is likely to have reached the target distribution.

3.3. Model comparison and selection

Watanabe-Alkaike Information Criterion (WAIC) (Watanabe 2010), a recent model comparison tool to measure the predictive accuracy of the fitted models in the Bayesian framework, will be used to compare the models. WAIC can be viewed as an improvement of Deviance Information Criterion(DIC) which, though popular, is known to be have some problems (Plummer 2008). WAIC is a fully Bayesian tool, closely approximates the Bayesian cross-validation, is invariant to reparameterisation and can be used for simple as well as hierarchical and mixture models.

4. Datasets

4.1. Telomerase data

Glas *et al.* (2003) systematically reviewed the sensitivity and specificity of cytology and other markers including telomerase for primary diagnosis of bladder cancer. They fitted a bivariate normal distribution to the logit transformed sensitivity and specificity values across the studies allowing for heterogeneity between the studies. From the included 10 studies, they reported that telomerase had a sensitivity and specificity of 0.75 [0.66, 0.74] and 0.86 [0.71, 0.94] respectively. They concluded that telomerase was not sensitive enough to be recommended for daily use. This dataset is available within the package and the following commands

```
R> data(telomerase)
```

```
R> telomerase
```

loads the data into the R enviroment and generates the following output

```
   ID Dis TP NonDis  TN
1   1  33 25     26  25
```

```
2   2  21 17     14  11
3   3 104 88     47  31
4   4  26 16     83  80
5   5  57 40    138 137
6   6  47 38     30  24
7   7  42 23     12  12
8   8  33 27     20  18
9   9  17 14     32  29
10 10  44 37     29   7
```

`ID` is the study identifier, `DIS` is the number of diseased, `TP` is the number of true positives, `NonDis` is the number of healthy and `TN` is the number of true negatives.

4.2. ASCUS triage data

Arbyn *et al.* (2013) performed a Cochrane review on the accuracy of human papillomavirus testing and repeat cytology to triage of women with an equivocal Pap smear to diagnose cervical precancer. They fitted the BRMA model in SAS using **METADAS** on 10 studies where both tests were used. They reported absolute sensitivity of 0.909 [0.857, 0.944] and 0.715 [0.629, 0.788] for HC2 and repeat cytology respectively. The specificity was 0.607 [0.539, 0.68] and 0.684 [0.599, 0.758] for HC2 and repeat cytology respectively. These data is used to demonstrate how the intercept only model is extended in a meta-regression setting. This dataset is also available within the package and the following commands

```
R> data(ascus)
```

```
R> ascus
```

loads the data into the R enviroment and generates the following output

```
   Test         StudyID  TP   FP  TN FN
1  RepC  Andersson 2005   6   14  28  4
2  RepC   Bergeron 2000   8   28  71  4
3  RepC Del Mistro 2010  20  191 483  7
4  RepC Kulasingam 2002  20   74 170  6
5  RepC     Lytwyn 2000   4   20  26  2
6  RepC      Manos 1999  48  324 570 15
7  RepC  Monsonego 2008  10   18 168 15
8  RepC      Morin 2001  14  126 214  5
9  RepC  Silverloo 2009  24   43 105 10
10 RepC    Solomon 2001 227 1132 914 40
11  HC2  Andersson 2005   6   17  25  4
12  HC2   Bergeron 2000  10   38  61  2
13  HC2 Del Mistro 2010  27  154 566  2
14  HC2 Kulasingam 2002  23  115 129  3
15  HC2     Lytwyn 2000   4   19  33  1
16  HC2      Manos 1999  58  326 582  7
```

```
17  HC2  Monsonego 2008  22  110  72  2
18  HC2       Morin 2001  17   88 253  2
19  HC2  Silverloo 2009  34   65  81  2
20  HC2     Solomon 2001 256 1050 984 11
```

`Test` is an explanatory variable showing the type of triage test, `StudyID` is the study identifier, `TP` is the number of true positives, `FP` is the number of false positives, `TN` is the number of true negatives, FN is the number of false negatives.

5. The intercept only model

The **CopulaDTA** package has five different correlation structures that result to five different bivariate beta-binomial distributions to fit to the data. The correlation structure is specified by indicating `copula ~=~"gauss"` or `"fgm"` or `"c90"` or `"270"` or `"frank"` in the `fitcopula` function. The Gaussian copula bivariate beta-binomial distribution is fitted to the `telomerase` data with the following code

```
R> gauss.1 <- cdtamodel("gauss")

R> fitgauss.1 <- fit(
+                 gauss.1,
+                 data = telomerase,
+                 SID = "ID",
+                 iter = 28000,
+                 warmup = 1000,
+                 thin = 30,
+                 seed = 3)
```

By default, `chains = 3` and `cores = 3` and need not be specified unless otherwise. From the code above, `28000` samples are drawn from each of the `3` chains, the first `1000` samples are discarded and thereafter every 30^{th} draw kept such that each chain has 900 post-warm-up draws making a total of 2700 post-warm-up draws. The seed value, `seed = 3`, specifies a random number generator to allow reproducibility of the results and `cores = 3` allows for parallel-processing of the chains by using `3` cores, one core for each chain. They were no initial values specified and in that case, the program randomly generates random values satisfying the parameter constraints. The trace plots in the top-left panel of Figure 1 produced with the code below show satisfactory mixing of the chains and convergence.

```
R> traceplot(fitgauss.1)
```

Next, obtain the model summary estimates as follows

```
R> print(fitgauss.1, digits = 4)

Posterior marginal mean sensitivity and specificity
with 95% credible intervals
```

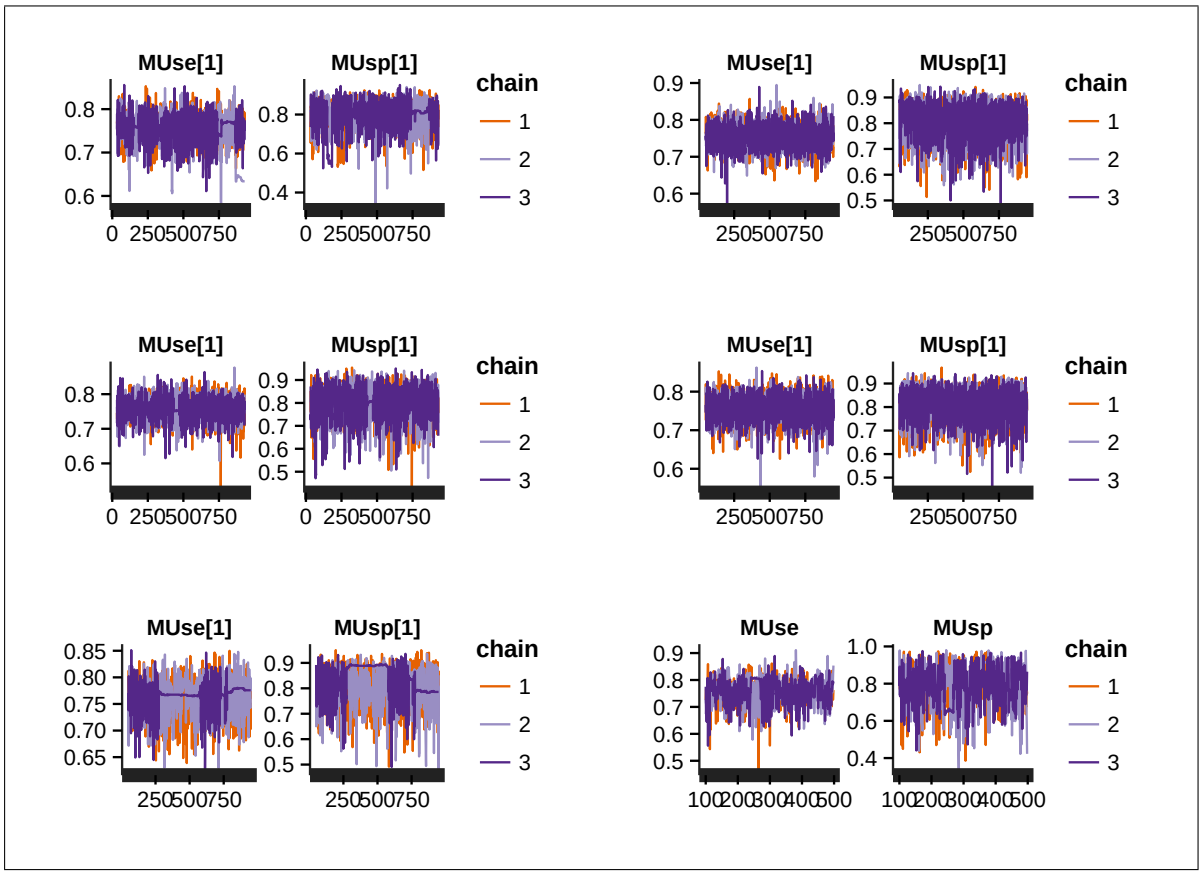

Figure 1: Trace plots of the posterior mean sensitivity and specificity for the `telomerase` data as estimated by the Gaussian, Clayton 90° (C90) and 270° (C270), Farlie-Gumbel-Morgenstern (FGM) and Frank copula based bivariate beta and bivariate normal (BRMA) distributions.

```
        Parameter    Mean   Lower   Upper n_eff  Rhat
MUse[1] Sensitivity  0.7540  0.6460  0.8119 122.4 1.010
MUsp[1] Specificity  0.8006  0.6235  0.9053 271.5 1.006
ktau[1] Correlation -0.8436 -0.9772 -0.3394 370.4 1.008

Model characteristics

Copula function: gauss, sampling algorithm: NUTS(diag_e)

Formula(1):  MUse ~ 1
Formula(2):  MUsp ~ 1
Formula(3):  Omega ~ 1
3 chain(s)each with iter=28000; warm-up=1000; thin=30.
post-warmup draws per chain=900;total post-warmup draws=2700.

Predictive accuracy of the model
```

```
Log point-wise predictive density (LPPD): -37.8529
Effective number of parameters: 7.0941
Watanabe-Akaike information Criterion (WAIC): 89.8940
```

From the output above, `n_eff` and `Rhat` both confirm proper mixing of the chains with little autocorrelation. The meta-analytic sensitivity `MUse[1]` and specificity `MUsp[1]` is 0.7540 [0.6460, 0.8119] and 0.8006 [0.6235, 0.9053] respectively. The Kendall's tau correlation between sensitivity and specificity is estimated to be -0.8436 [-0.9772, -0.3394].

The command below produces a forest plot in Figure 2.

```
R> plot(model1, graph = 3, title.3 = "" )
```

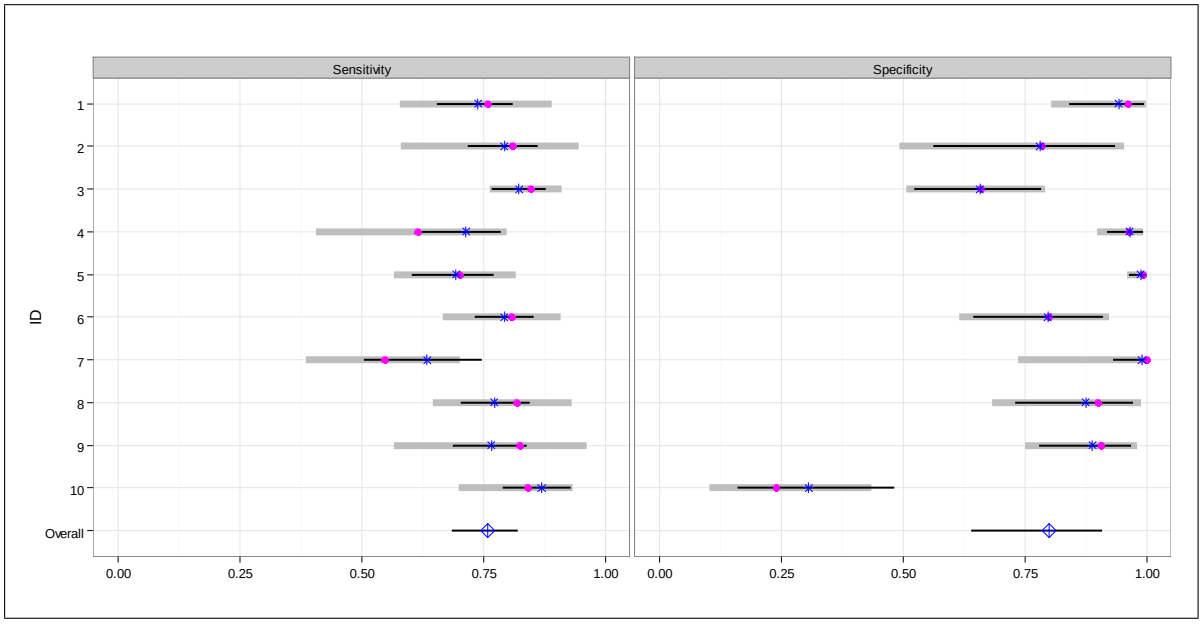

Figure 2: Plot of the study-specific sensitivity and specificity (magenta points) and their corresponding 95 % exact confidence intervals (thick grey lines), superimposed with the posterior estimates (blues stars) and their corresponding 95 % credible intervals (think black lines). Posterior estimates from the Gaussian copula based bivariate beta distribution for the `telomerase` data.

As observed in Figure 2, the posterior study-specific sensitivity and specificity are less extreme and variable than the 'observed' study-specific sensitivity and specificity. In other words, there is 'shrinkage' towards the overall mean sensitivity and specificity as studies borrow strength from each other in the following manner: the posterior study-specific estimates depends on the global estimate and thus also on all other the studies.

The mean sensitivity and specificity as estimated by the other four copula based bivariate beta distributions are in Table 1 and graphically shown in Figure 3. Though not presented here, the full code of the other four fitted copula based bivariate beta distributions is in the replication code. Figure 1 shows satisfactory chain mixing with little autocorrelation

apart from the 'Clayton270' model. The Clayton copula is known to be unstable when the correlation parameter is close to the boundaries (-1 or 0) and this could be the reason why sampling from the posterior distribution was difficult.

For comparison purpose, the current recommended model; the BRMA, which uses normal marginals is also fitted to the data though it is not part of the **CopulaDTA** package. The model is first expressed in Stan modelling language in the code below and is stored within R environment as character string named `BRMA1`.

```
R> BRMA1 <- "
data{
        int<lower = 0> Ns;
        int<lower = 0> tp[Ns];
        int<lower = 0>  dis[Ns];
        int<lower = 0>  tn[Ns];
        int<lower = 0>  nondis[Ns];
}
parameters{
        real etarho;
        vector[2] mul;
        vector<lower = 0>[2] sigma;
        vector[2] logitp[Ns];
        vector[2] logitphat[Ns];
}
transformed parameters{
        vector[Ns] p[2];
        vector[Ns] phat[2];

        real MU[2];
        vector[2] mu;
        real rho;
        real ktau;
        matrix[2,2] Sigma;

        rho <- tanh(etarho);
        ktau <- (2/pi())*asin(rho);

        for (a in 1:2){
                for (b in 1:Ns){
                        p[a][b] <- inv_logit(logitp[b][a]);
                        phat[a][b] <- inv_logit(logitphat[b][a]);

                }
                mu[a] <- inv_logit(mul[a]);
        }

        MU[1] <- mean(phat[1]);
        MU[2] <- mean(phat[2]);
```

```
        Sigma[1, 1] <- sigma[1]^2;
        Sigma[1, 2] <- sigma[1]*sigma[2]*rho;
        Sigma[2, 1] <- sigma[1]*sigma[2]*rho;
        Sigma[2, 2] <- sigma[2]^2;
}
model{
        etarho ~ normal(0, 10);
        mul ~ normal(0, 10);
        sigma ~ cauchy(0, 2.5);

        for (i in 1:Ns){
                logitp[i] ~ multi_normal(mul, Sigma);
                logitphat[i] ~ multi_normal(mul, Sigma);
        }

        tp ~ binomial(dis,p[1]);
        tn ~ binomial(nondis, p[2]);

}
generated quantities{
        vector[Ns*2] loglik;

        for (i in 1:Ns){
                loglik[i] <- binomial_log(tp[i], dis[i], p[1][i]);
        }
        for (i in (Ns+1):(2*Ns)){
                loglik[i] <- binomial_log(tn[i-Ns], nondis[i-Ns], p[2][i-Ns]);
        }
}
"
```

Next, prepare the data by creating as list as follows

```
R> datalist = list(
+         tp = telomerase$TP,
+         dis = telomerase$TP + telomerase$FN,
+         tn = telomerase$TN,
+         nondis = telomerase$TN + telomerase$FP,
+         Ns = 10)
```

In the `data` block the dimensions and names of variables in the dataset are specified, here `Ns` indicate the number of studies in the dataset. The `parameters` block introduces the unknown parameters to be estimated. These are `etarho`; a scalar representing the Fisher's transformed form of the association parameter ρ, `mul`;a 2×1 vector representing the mean of sensitivity and specificity on the logit scale for a central study where the random-effect is zero, `sigma`; a 2×1 vector representing the between study standard deviation of sensitivity and specificity

on the logit scale, `logitp`; a *Ns* × *2* array of study-specific sensitivity in the first column and specificity in the second column on logit scale, and `logitphat`; a *Ns* × *2* array of predicted sensitivity in the first column and predicted specificity in the second column on logit scale.

The parameters are further transformed in the `transformed parameters` block. Here, `p` is a *2* × *Ns* array of sensitivity in the first column and specificity in the second column after inverse logit transformation of `logitp`, and `phat` is a *2* × *Ns* array of predicted sensitivity in the first column and predicted specificity in the second column after inverse logit transformation of `logitphat` to be used in computing the meta-analytic sensitivity and specificity. `mu` is a *2* × *1* vector representing the mean of sensitivity and specificity for a certain study with a random effect equal to 0, `MU` is a *2* × *1* vector containing the meta-analytic sensitivity and specificity, `Sigma`; a *2* × *2* matrix representing the variance-covarince matrix of sensitivity and specificity on the logit scale, `rho` and `ktau` are scalars representing the Pearson's and Kendall's tau correlation respectively. The prior distributions for the all parameters and data likelihood are defined in the `model` block. Finally, in the `generated quantities` block, `loglik` is a *(2Ns)* × *1* vector of the log likelihood needed to compute the WAIC.

Next, call the function `stan` from the **rstan** package to translate the code into C++, compile the code and draw samples from the posterior distribution as follows

```
R> brma.1 <- stan(model_code = BRMA1,
+                 data = datalist,
+                 chains = 3,
+                 iter = 5000,
+                 warmup = 1000,
+                 thin = 10,
+                 seed = 3,
+                 cores = 3)
```

The parameter estimates are extracted and the chain convergence and autocorrelation examined further with the following code

```
R> print(brma.1, pars = c('MU', 'mu', 'rho'),
+        digits = 4,
+        prob=c(0.025, 0.975))
```

The above lines of code print the following output

```
Inference for Stan model: d6a1713b5981968874c97152db2bb815.
3 chains, each with iter=5000; warmup=1000; thin=10;
post-warmup draws per chain=400, total post-warmup draws=1200.

                mean se_mean     sd    2.5%   97.5% n_eff   Rhat
MU[1]         0.7549  0.0020 0.0490  0.6438  0.8408   594 0.9995
MU[2]         0.7901  0.0056 0.1121  0.5252  0.9554   397 1.0036
mu[1]         0.7681  0.0014 0.0367  0.6850  0.8425   714 0.9999
mu[2]         0.8971  0.0026 0.0715  0.7189  0.9822   763 1.0058
rho          -0.9338  0.0092 0.1224 -0.9993 -0.5711   175 1.0208
```

```
Samples were drawn using NUTS(diag_e) at Mon Dec 07 14:59:26 2015.
For each parameter, n_eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor on split chains (at
convergence, Rhat=1).
```

The meta-analytic sensitivity (`MU[1]`) and specificity (`MU[2]`) and 95% credible intervals are 0.7549[0.6438, 0.8408] and 0.7901[0.5252, 0.9554] respectively. This differs from what the authors published (0.75[0.66, 0.74] and 0.86[0.71, 0.94]) in two ways. The authors fitted the standard bivariate normal distribution to the logit transformed sensitivity and specificity values across the studies allowing for heterogeneity between the studies as expressed in Equation 6 and disregarded the higher level of the hierarchical model expressed in Equation 5. Because of this the authors had to use a continuity correction of 0.5 since the seventh study had 'observed' specificity equal to 1, a problem not encountered in the hierarchical model. Secondly the authors do not report the meta-analytic values but rather report the mean sensitivity(`mu[1]`) and specificity (`mu[2]`) for a particular, hypothetical study with random-effect equal to zero, which in our case is 0.7681[0.6850, 0.8425] and 0.8971[0.7189, 0.98227] respectively and is comparable to what the authors reported. This discrepancy between `MU` and `mu` will indeed increase with increase in the between study variability.

Model	Parameter	Mean	Lower	Upper	n_eff	Rhat	WAIC
Gaussian	Sensitivity	0.7540	0.6460	0.8119	122	1.0076	89.8940
	Specificity	0.8006	0.6235	0.9053	271	1.0061	
	Correlation	-0.8436	-0.9772	-0.3394	370	1.0083	
C90	Sensitivity	0.7579	0.6894	0.8165	2046	1.0008	92.4859
	Specificity	0.7996	0.6352	0.9100	1473	1.0053	
	Correlation	-0.7338	-0.9830	0.0000	1347	1.0000	
C270	Sensitivity	0.7606	0.6888	0.8165	807	1.0097	90.5935
	Specificity	0.8125	0.6439	0.9048	61	1.0311	
	Correlation	-0.7526	-0.9800	0.0000	17	1.0601	
FGM	Sensitivity	0.7576	0.6901	0.8169	2700	0.9994	95.4854
	Specificity	0.8053	0.6427	0.9085	2700	1.0001	
	Correlation	-0.1871	-0.2222	0.2222	2377	1.0012	
Frank	Sensitivity	0.7576	0.6880	0.8165	2693	1.0008	89.9743
	Specificity	0.8097	0.6526	0.9112	2582	1.0007	
	Correlation	-0.7073	-0.8504	-0.1852	2700	NA	
BRMA	Sensitivity	0.7549	0.6438	0.8408	594	0.9995	86.6359
	Specificity	0.7901	0.5252	0.9554	397	1.0036	
	Correlation	-0.8204	-0.9755	-0.3870	39	1.0556	

Table 1: The posterior mean, 95% credible interval, effective sample size and potential scale reduction factor $\hat{R}$ factor for the marginal means and correlation parameters as estimated by the Gaussian, Clayton 90° (C90) and 270° (C270), Farlie-Gumbel-Morgenstern (FGM) and Frank copula based bivariate beta and bivariate normal (BRMA) distributions for the `telomerase` data.

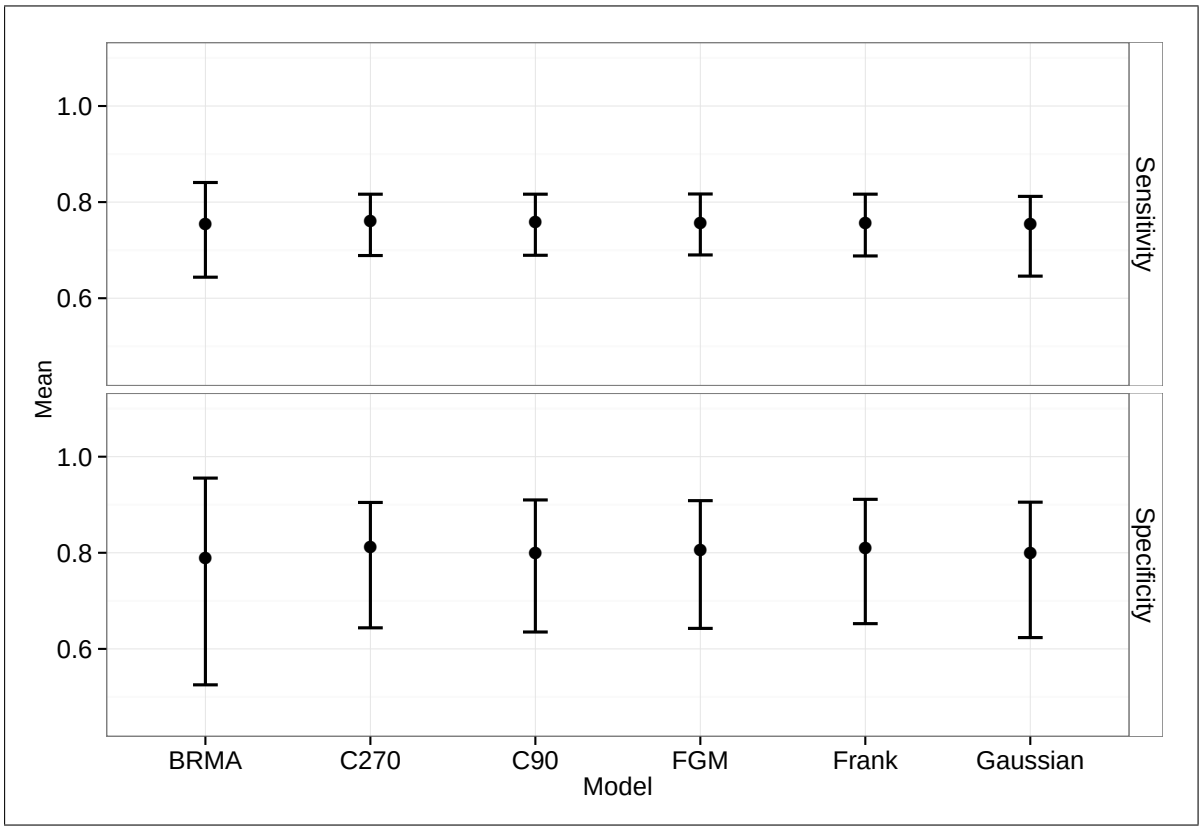

Figure 3: Plot of the posterior meta-analytic sensitivity (upper) and specificity (lower) and the correspondinb 95% credible intervals) as estimated by the Gaussian, Clayton 90° (C90) and 270° (C270), Farlie-Gumbel-Morgenstern (FGM) and Frank copula based bivariate beta and bivariate normal (BRMA) distributions for the `telomerase` data.

5.1. Model comparison

Table 1 shows that the correlation as estimated by the BRMA model and the Gaussian copula bivariate beta are more extreme but comparable to the estimates from the Frank, 90°- and 270°- Clayton copula. On the other extreme is the estimate from the model FGM copula bivariate beta and this is due to the constraints on the association parameter in the FGM copula where values lie within |2/9|.

In Figure 3, the marginal mean sensitivity and specificity from the five bivariate beta distributions are comparable with subtle differences in the 95 percent credible intervals despite differences in the correlation structure.

Glas *et al.* (2003) and Riley *et al.* (2007a) estimated the Pearson's correlation parameter in the BRMA model ρ as -1 within the frequentist framework. Using maximum likelihood estimation, Riley *et al.* (2007b) showed that the between-study correlation from the BRMA is often estimated as +/-1. Without estimation difficulties, Table 1 shows an estimated Pearson's correlation of -0.9338[-0.9993, -0.57118]. This is because Bayesian methods are not influenced by sample size and therefore able to handle cases of small sample sizes with less issues.

Essentially, all the six models are equivalent in the first level of hierarchy and differ in spec-

ifying the prior distributions for the 'study-specific' sensitivity and specificity. As thus, the models should have the same number of parameters in which case it makes sense then to compare the log predictive densities. Upon inspection, the log predictive densities from the six models are practically equivalent (min=37.40, max=38.62) but the effective number of parameters differed a bit (max=5.9, max=9.1). Apparently, the last column of Table 1 indicates that the BRMA fits the data best based on the WAIC.

6. Meta-regression

The `ascus` dataset has `Test` as a covariate. The covariate is used as it is of interest to study its effect on the joint distribution of sensitivity and specificity (including the correlation). The following code fits the FGM copula based bivariate beta-binomial distribution to the data

```
R> fgm.2 <- cdtamodel(copula = "fgm",
+                     modelargs = list(formula.se = StudyID ~ Test + 0))

R> fitfgm.2 <- fit(fgm.2,
+                  data = ascus,
+                  SID = "StudyID",
+                  iter = 19000,
+                  warmup = 1000,
+                  thin = 20,
+                  seed = 3)
```

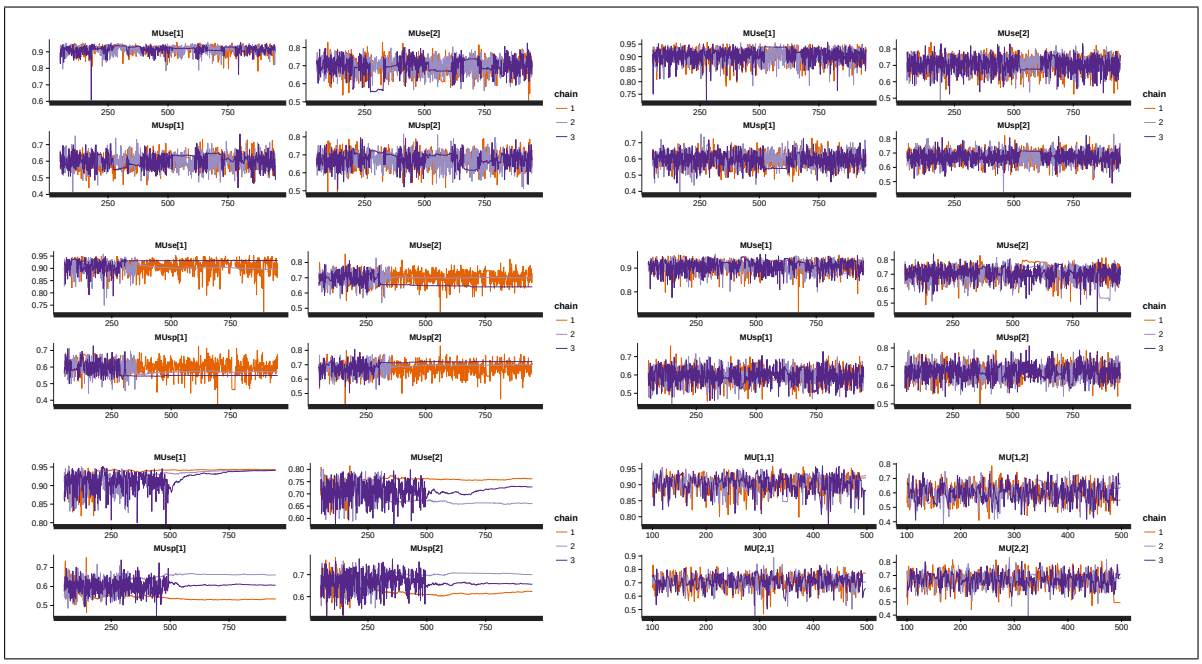

Figure 4: Trace plots of the posterior mean sensitivities and specificities for the `ascus` data as estimated by the Gaussian, Clayton 90° (C90) and 270° (C270), Farlie-Gumbel-Morgenstern (FGM) and Frank copula based bivariate beta and bivariate normal (BRMA) distributions.

Figure 4 shows the trace plots for all the six models fitted to the `ascus` data where all parameters, including the correlation parameter(except the BRMA) are modeled as a function of the covariate. There is proper chains mixing and convergence except for the case of the Clayton copula based bivariate beta. From the posterior relative sensitivity and specificity plotted in Figure 5, all the models that converged generally agree that repeat cytology was less sensitive than HC2 without significant loss in specificity.

Model	Test	Mean	Lower	Upper	n_eff	Rhat	WAIC
Gaussian	HC2	-0.4799	-0.9902	0.8972	403	1.0027	5350.6348
	Repc	-0.9164	-0.9972	-0.6122	390	1.0004	
C90	HC2	-0.1122	-0.9238	0.0000	14	1.0910	5347.9687
	Repc	-0.8554	-0.9825	-0.3763	38	1.0553	
C270	HC2	-0.0524	-0.8139	0.0000	43	1.0926	5339.1644
	Repc	-0.7898	-0.9783	-0.3691	5	1.6967	
FGM	HC2	-0.0836	-0.2222	0.2222	451	1.0072	5355.6485
	Repc	-0.1999	-0.2222	0.1704	2584	0.9994	
Frank	HC2	-0.5238	-0.8201	0.6067	2700	NA	5353.1396
	Repc	-0.7410	-0.8627	-0.3253	2700	NA	
BRMA	Both	-0.8483	-0.9954	-0.4391	5	1.1967	5348.3116

Table 2: The posterior mean, 95% credible intervals, effective sample size, potential scale reduction factor of the correlation parameter(s) as estimated by the Gaussian, Clayton 90° (C90) and 270° (C270), Farlie-Gumbel-Morgenstern (FGM) and Frank copula based bivariate beta and bivariate normal (BRMA) distributions for the `ascus` dataset.

The `n_eff` in Table 2 indicate substantial autocorrelation in sampling the correlation parameters except in the 'Gaussian', 'FGM' and 'Frank' models. From the copula based bivariate beta distributions, it is apparent that the correlation between sensitivity and specificity in HC2 and repeat cytology is different.

The 'Clayton' models have the lowest WAIC even though sampling from the posterior distribution was difficult as seen in their trace plots in Figure 4 and the `n_eff` and `Rhat` in Table 2. The difficulty in sampling from the posterior could be signalling over-parameterisation of the correlation structure. It would thus be interesting to re-fit the models using only one correlation parameter and compare the models. WAIC is known to fail in certain settings and this examples shows that it is crucial to check the adequacy of the fit and plausibility of the model and not blindly rely on an information criterion to select the best fit to the data.

7. Discussion

Copula-based models offer great flexibility and ease but their use is not without caution. While the copulas used in this paper are attractive as they are mathematically tractable, Mikosch (2006) and Genest and Remillard (2006) noted that it might be difficult to estimate copulas from data. Furthermore, the concepts behind copula models is slightly more complex and therefore require statistical expertise to understand and program them as they are not yet available as standard procedure/programs in statistical software.

In this paper, several advanced statistical models for meta-analysis of diagnostic accuracy

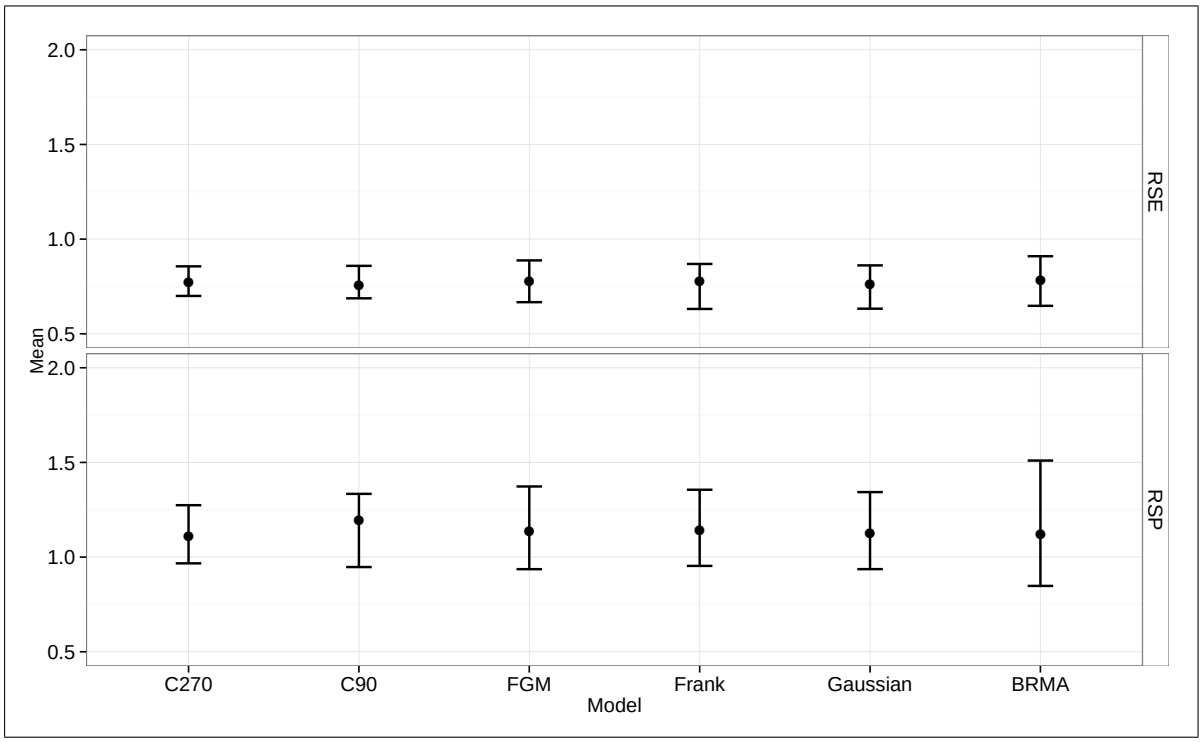

Figure 5: Pooled relative sensitivity (on top) and relative specificity (bottom) of repeat cytology (posterior mean and 95% credible intervals) compared to HPV testing with HC2 to detect cervical precancer in women with an atypical Pap smear estimated.

studies were briefly discussed. The use of the R package **CopulaDTA** within the flexible **Stan** interface was demonstrated and shows how complex models can be implemented in a convenient way.

In most practical situations, the marginal mean structure is of primary interest and the correlation structure is treated a nuisance making the choice of copula less critical. Nonetheless, an appropriate correlation structure is critical in the interpretation of the random variation in the data as well as obtaining valid model-based inference for the mean structure.

When the model for the mean is correct but the true distribution is misspecified, the estimates of the model parameters will be consistent but the standard errors will be incorrect Agresti (2002). Nonetheless, the bivariate beta distribution has the advantage to allow direct joint modelling of sensitivity and specificity, without the need of any transformation, and consequently providing estimates with the appropriate meta-analytic interpretation but with the disadvantage of being more computationally intensive for some of the copula functions.

Leeflang *et al.* (2013) showed that the sensitivity and specificity often vary with disease prevalence. The models presented above can easily be extended and implemented to jointly model prevalence, sensitivity and specificity using tri-variate copulas.

There were some differences between the models in estimating the meta-analytic sensitivity and specificity and the correlation. Therefore, further research is necessary to investigate the effect of certain parameters, such as the number of studies, sample sizes and misspecification of the joint distribution on the meta-analytic estimates.

8. Conclusion

The proposed Bayesian joint model using copulas to construct bivariate beta distributions, provides estimates with both the appropriate marginal as well as conditional interpretation, as opposed to the typical BRMA model which estimates sensitivity and specificity for specific studies with a particular value for the random-effects. Furthermore, the models do not have estimation difficulties with small sample sizes or large between-study variance because: **i**) the between-study variances are not constant but depends on the underlying means and **ii** Bayesian methods are less influenced by small samples sizes.

The fitted models generally agree that the mean specificity was slightly lower than what Glas *et al.* (2003) reported and based on this we conclude that telomerase was not sensitive and specific enough to be recommended for daily use.

In the ASCUS triage data, conclusion based on the fitted models is in line with what the authors conclude: that HC2 was considerably more sensitive but sligthly and non-significantly less specific than repeat cytology to triage women with an equivocal Pap smear to diagnose cervical precancer.

While the BRMA had the lowest WAIC for both datasets, we still recommend modelling of sensitivity and specificity using bivariate beta distributions as they easily and directly provide meta-analytic estimates.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

M. Arbyn designed the OPSADAC project (Optimisation of statistical procedures to assess the diagnostic accuracy of cervical cancer screening tests) of which this study is a part of. Victoria and M. Aerts conceptualized and initiated the study. Victoria wrote the code, analysed the data and drafted manuscript. M. Arbyn and M. Aerts edited the manuscript. All authors reviewed and approved the final manuscript.

Acknowledgements

V. Nyaga recieved financial support from the Scientific Institute of Public Health (Brussels) through the OPSADAC project. M. Arbyn was supported by the COHEAHR project funded by the 7th Framework Programme of the European Commission (grant No 603019). M. Aerts was supported by the IAP research network nr P7/06 of the Belgian Government (Belgian Science Policy).

References

Agresti A (2002). *Categorical Data Analysis.* 2nd edition. John Wiley & Sons, New-York, 131.

Arbyn M, *et al.* (2013). "Human Papillomavirus Testing Versus Repeat Cytology for Triage of Minor Cytological Cervical Lesions." *Cochrane Database of Systematic Reviews*, 31–201.

Arends LR, *et al.* (2008). "Bivariate Random Effects Meta-analysis of ROC Curves." *Medical Decision Making*, **28**(5), 621–638.

Arnold BC, Ng HKT (2011). "Flexible Bivariate Beta Distributions." *Journal of Multivariate Analysis*, **102**(8), 1194–1202.

Carpenter B, *et al.* (In press). "Stan: A Probabilistic Programming Language." *Journal of Statistical Software.*

Chu H, Cole S (2006). "Bivariate Meta-Analysis of Sensitivity and Specificity with Sparse Data: A Generalized Linear Mixed Model Approach." *Journal of Clinical Epidemiology*, **59**(12), 1331–1332.

Clayton DG (1978). "A model for Association in Bivariate Life Tables and its Application in Epidemiological Studies of Familial Tendency in Chronic Disease Incidence." *Biometrika*, **65**(1), 141–151.

Cong X, Cox DD, Cantor SB (2007). "Bayesian Meta-analysis of Papanicolaou Smear Accuracy." *Gynecologic Oncology*, **107**(1 Suppl 1), S133–S137.

Farlie DGJ (1960). "The Performance of Some Correlation Coefficients for a General Bivariate Distribution." *Biometrika*, **47**, 307–323.

Frank MJ (1979). "On the Simultaneous Associativity of F(x, y) and x + y - F(x, y)." *Aequationes Mathematicae*, 194–226.

Genest C, Remillard B (2006). "Comments on T. Mikosh's paper, Copulas: Tales and Facts." *Extremes*, **9**, 27–36.

Glas AS, *et al.* (2003). "Tumor Markers in the Diagnosis of Primary Bladder Cancer. A Systematic Review." *The Journal of Urology*, **169**(6), 1975–1982.

Gumbel EJ (1960). "Bivariate Exponential Distributions." *Journal of the American Statistical Association*, **55**, 698–707.

Hoffman MD, Gelman A (In press). "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research.*

Leeflang MM, *et al.* (2013). "Variation of a Test's Sensitivity and Specificity with Disease Prevalence." *Canadian Medical Association Journal*, **185**(11), E537–E544.

Libby DL, Novick RE (1982). "Multivariate Generalized Beta Distributions with Applications to Utility Assessment." *Journal of Educational Statistics*, **7**(4), 271–294.

Meyer C (2013). "The Bivariate Normal Copula." *Communications in Statistics - Theory and Methods*, **42**(13), 2402–2422.

Mikosch TV (2006). "Copulas: Tales and Facts. Discussion Paper with a Rejoinder." *Extremes*, **9**, 3–20, 55–62.

Molenberghs G and Verbeke G. (2005). *Models for Discrete Longitudinal Data.* Springer-Verlag, New York, 259, 267.

Morgenstern D (1956). "Einfache Beispiele Zweidimensionaler Verteilungen." *Mitteilungsblatt fürMathematische Statistik*, **8**, 234–235.

Moses LE, Shapiro D and Littenberg B (1993). "Combining Independent Studies of a Diagnostic Test into a Summary ROC Curve: Data-Analytic Approaches and some Additional Considerations." *Statistics in Medicine*, **12**(14), 1293–1316.

Nelsen RB (2006). *An Introduction to Copulas*. Springer-Verlag, New York.

Nikoloulopoulos AK (2015a). "A Mixed Effect Model for Bivariate Meta-analysis of Diagnostic Test Accuracy Studies using a Copula Representation of the Random Effects Distribution." *Statistics in Medicine*, **34**(29), 3842–3865.

Nikoloulopoulos AK (2015b). ***CopulaREMADA**: Copula Random Effects Model for Bivariate and Trivariate Meta-Analysis of Diagnostic Test Accuracy Studies*. R package version 0.9, URL https://CRAN.R-project.org/package=CopulaREMADA.

Nyaga VN (2015). ***CopulaDTA**: Copula Based Bivariate Beta-Binomial Model for Diagnostic Test Accuracy Studies*. R package version 0.0.2, URL http://CRAN.R-project.org/package=CopulaDTA.

Olkin I, Liu R (2003). "A Bivariate Beta Distribution." *Statistics and Probability Letters*, **62**, 407–412.

Olkin I, Trikalinos TA (2014). "Constructions for a Bivariate Beta Distribution." *ARXIV*, 1–10.

Patton AJ (2006). "Modelling Asymmetric Exchange Rate Dependence." *International Economic Review*, **47**(2), 527–556.

Plummer M (2003). "JAGS: A program for Analysis of Bayesian Graphical Models using Gibbs Sampling." *Proceedings of the 3rd international Workshop on Distributed Statistical Computing*, Technische Universit at Wien, **124**.

Plummer M (2008). "Penalized Loss Functions for Bayesian Model Comparison." *Biostatistics*, **9**, 523–539.

R Core Team (2015). *R : A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

Reitsma JB, Glas AS, Rutjes AW, Scholten RJ, Bossuyt PM, Zwinderman AH (2005). "Bivariate Analysis of Sensitivity And Specificity Produces Informative Summary Measures in Diagnostic Reviews." *Journal of Clinical Epidemiology*, **58**(10), 982–990.

Riley RD (2009). "Multivariate Meta-analysis: The Effect of Ignoring Within-study Correlation." *Journal of the Royal Statistical Society*, **172**(4).

Riley RD, Abrams KR, Lambert PC, Sutton AJ, Thompson JR (2007a). "An Evaluation of Bivariate Random-effects Meta-analysis for the Joint Synthesis of Two Correlated Outcomes." *Statistics in Medicine*, **26**(1), 78–97.

Riley RD, Abrams KR, Sutton AJ, Lambert PC, Thompson JR (2007b). "Bivariate Random-Effects Meta-Analysis and the Estimation of Between-Study Correlation." *BMC Medical Research Methodology*, **7**(3).

Rutter CM, Gatsonis CM (2001). "A Hierarchical Regression Approach to Meta-Analysis of Diagnostic Test Accuracy Evaluations." *Statistics in Medicine*, **20**, 2865–84.

Sarmanov O (1966). "Generalized Normal Correlation and Two-Dimensional Fréchet Classes." *Soviet Mathematics - Doklady*, **7**, 596–599.

Sklar A (1959). "Fonctions de Répartition á n Dimensions et Leurs Marges." *Publications de l'Institut de Statistique de L'Université de Paris*, **8**, 229–231.

Stan Development Team (2015). *Stan : A C++ Library for Probability and Sampling.* Version 2.8.0. URL http://mc-stan.org/.

Takwoingi Y, Guo B, Riley RD and Deeks, JJ (2015). "Performance of Methods for Meta-Analysis of Diagnostic Test Accuracy with Few Studies or Sparse Data." *Statistical Methods in Medical Research*, **0**(0), 1–19.

Vehtari A, Gelman A (2014). "WAIC and Cross-Validation in Stan". *Unpublished*.

Watanabe S (2010). "Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory." *Journal of Machine Learning Research*, **11**, 3571–3594.

Affiliation:

Victoria N Nyaga
Unit of Cancer Epidemiology
Scientific Institute of Public Health
Juliette Wytsmanstraat 14
1050 Brussels, Belgium
E-mail: victoria.nyawiranyaga@wiv-isp.be
E-mail: victoria.nyaga@uhasselt.be

Marc Arbyn
Unit of Cancer Epidemiology
Scientific Institute of Public Health
Juliette Wytsmanstraat 14
1050 Brussels, Belgium
E-mail: marc.arbyn@wiv-isp.be

Marc Aerts
Center for Statistics
Hasselt University
Martelarenlaan 42
3500 Hasselt, Belgium
E-mail: marc.aerts@uhasselt.be